

Solution Management



SAP Enterprise Portal 6.0



Copyright

© Copyright 2002 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft®, WINDOWS®, NT®, EXCEL®, Word®, PowerPoint® and SQL Server® are registered trademarks of Microsoft Corporation.

IBM®, DB2®, DB2 Universal Database, OS/2®, Parallel Sysplex®, MVS/ESA, AIX®, S/390®, AS/400®, OS/390®, OS/400®, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere®, Netfinity®, Tivoli®, Informix and Informix® Dynamic Server™ are trademarks of IBM Corporation in USA and/or other countries.

ORACLE® is a registered trademark of ORACLE Corporation.

UNIX®, X/Open®, OSF/1®, and Motif® are registered trademarks of the Open Group.

Citrix®, the Citrix logo, ICA®, Program Neighborhood®, MetaFrame®, WinFrame®, VideoFrame®, MultiWin® and other Citrix product names referenced herein are trademarks of Citrix Systems, Inc.

HTML, DHTML, XML, XHTML are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

JAVA® is a registered trademark of Sun Microsystems, Inc.

JAVASCRIPT® is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, SAP Logo, R/2, RIVA, R/3, SAP ArchiveLink, SAP Business Workflow, WebFlow, SAP EarlyWatch, BAPI, SAPHIRE, Management Cockpit, mySAP, mySAP.com, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. MarketSet and Enterprise Buyer are jointly owned trademarks of SAP Markets and Commerce One. All other product and service names mentioned are the trademarks of their respective owners.

Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

Typographic Conventions

Type Style	Description
<i>Example text</i>	Words or characters that appear on the screen. These include field names, screen titles, pushbuttons as well as menu names, paths and options. Cross-references to other documentation.
Example text	Emphasized words or phrases in body text, titles of graphics and tables.
EXAMPLE TEXT	Names of elements in the system. These include report names, program names, transaction codes, table names, and individual key words of a programming language, when surrounded by body text, for example, SELECT and INCLUDE.
Example text	Screen output. This includes file and directory names and their paths, messages, source code, names of variables and parameters as well as names of installation, upgrade and database tools.
EXAMPLE TEXT	Keys on the keyboard, for example, function keys (such as F2) or the ENTER key.
Example text	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example text>	Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries.

Contents

Solution Management	5
Backup and Restore	6
Backup Strategy	7
Backup Types	8
SAP Enterprise Portal 6.0	10
Identifying Critical Components	12
Enterprise Portal Components	14
Windows	15
Unix Platform	15
Portal Server	16
Knowledge Management (KM)	17
Text Retrieval and Information Extraction (TREX) Software Backup	18
Content Management (CM)	19
User Management Engine (UME)	20
Restoring	22
Post - Recovery Checks	24



Solution Management

The task of implementing, and keeping your SAP Enterprise Portal up and running round the clock at peak performance has never been more vital to your business success.

This documentation provides a central starting point for the solution management to keep your solution up and running. Currently, it contains information for implementing a backup and restore strategy for SAP Enterprise Portal.



Backup and Restore

Backup refers to the activity of copying files, data, or network volume, with the intention of preserving them for later use in case of hardware failure, or other disaster. When you retrieve files that have been backed up earlier, you are restoring them.

Purpose

There are many reasons for backing up data such as, hard drives crash, viruses infecting your system and destroying your data, human error resulting in havoc on your business, or a stolen system with your data on it.

Therefore, it is important to implement a backup and restore strategy that protects your system against data loss, and enables you to restore the system to its correct and consistent state. The most important part of any security strategy is to backup the system at regular intervals. This means the system (files, databases, etc.) must be copied to another storage medium. When the system is damaged, the stored duplicate can be re-loaded to restore it.

Backup effectively contributes to safeguarding your system if it is performed as part of an overall backup and restore strategy. It is therefore important to first, carefully plan, prepare, and then test the backup and restore system before applying it in a productive system.

In planning for backup and restore, define details for performing backups, what should be backed up, and at what times, which media to use, and how to verify backups to make sure that you have meaningful stored duplicate data that can help you to restore your systems later. In addition, many different factors play a role when planning your strategy; for example, you need to take into account the transaction workload, the maximum permissible downtime, the available hardware, and the amount of data loss that is tolerable.



Backup Strategy

The purpose of backups is to ensure that the files and data copied as a backup to a system can be used to restore the same system (or another system) to its correct and consistent state after it has been damaged. An effective backup strategy must enable a safe and efficient recovery of a damaged system.

An effective planning strategy that provides adequate insurance against data loss is always based on appropriate hardware and system configuration, and describes a backup cycle of one month. Such planning usually includes the following:

- Identified parts/area (drives, directories, files) of the system to backup
- Suitable type of backup
- Type of recovery mode to set for the SAP system and databases
- The frequency with which you perform backups
- Appointed times for performing backups
- Type of storage media (devices) to use for backup
- Backup management, and maintenance issues
- Monitoring and verifying backups

This document does not provide detailed information on available backup strategies, moreover, information on different backup and restore strategies relate to the specific backup and restore solution offered by each vendor.



Backup Types

There are several ways to perform backups, including the following:

- Online and offline backups
- Media rotation methods

Online and Offline Backups

On-line backup takes place during system operations (while various components and services are running), as such it can be very inconsistent and needs additional steps to completely restore a damaged system.

Offline backup is performed after system operations have been stopped, and a snapshot of the system's state is mirrored onto a media at a point in time.

Besides, online backup capabilities of databases, and the file system have the disadvantage that opened files are not backed up. Therefore, special backup software is required to complement the method.

A backup strategy has to include offline and online backups for 7x24 hours support. The benefit of offline backups is the ability to preserve the consistent state of the backed up system, and then to restore the system to the same state afterwards.

Media Rotation Methods

Media rotation is the use of multiple media devices to perform backup. There are different rotation schemes suitable to different systems. Backup strategies for media rotation differ mostly by the number of media required and the scheduled retention period for each media.

Two common media rotation methods are:

- **Full Backup:** Full backup refers to the activity of backing up everything on your system. This backup method must be performed regularly, at least once a week, depending on your work volume.

Advantage of full backup is that files are easy to find, and you do not have to search through several media for requested files. Also, it provides the most current backup of your entire system, and easier to restore.

Redundant backup easily occurs, since most of the files on your file server rarely change. In addition, each full backup includes a copy of what has already been backed up.

Also, full backup requires more media, and takes longer to perform. Since it affects the speed of your system, you should perform this backup at a time when the workload is low.

- **Modified Backup:** Modified backup can be either incremental or differential.

Incremental backup is backing up only new files created, or those that have changed since the last backup was performed. To restore a system from an incremental backup, you need the last full backup and each incremental backup performed.

Advantages of incremental backup include, better usage of media, less time required for backup. The main disadvantage is that backups are spread across multiple media

Differential backup copies new file created, or those that have changed since the last backup. This is similar to incremental backup, however, to restore a system from a differential backup, you need the full backup and the last differential backup performed.

Some of the advantages of differential backup are that it takes less amount of time, fewer media devices are used, and the level and risk of media errors are minimal for backup and restore.

The following are some of the possible backup strategies you can combine:

- **Son:** - The son involves doing a full backup every day.
- **Father and Son:** - The father and son combination involves both full and differential, or incremental backup scheduled fortnightly.
- **Grandfather:** - The grandfather scheme uses several media devices over a period of time:
 - a. Four media devices are used from Monday - Thursday for incremental or differential backup.
 - b. A set of three media devices is used every Friday for full backup.
 - c. 12 media devices are used for monthly full backups and are kept off-site.

Before implementing a backup strategy, you need to decide whether you want to perform only full backup, or use a strategy that includes full backup and one of the modified types (differential or incremental). There are advantages and disadvantages to each backup type.



SAP Enterprise Portal 6.0

Backup and restore strategy for the SAP Enterprise Portal infrastructure must be carefully developed to achieve a consistent state for the system upon recovery. The strategy may involve backup at the component level, or a full backup of the entire system.

Enterprise Portal infrastructure for customers differ, therefore performing a full backup is a comprehensible solution. This solution can easily restore portal components and enable the portal to recover from a complete system loss.

In addition, portal components and their configurations are interdependent on each, and need to be backed up together.

Application Data (Business Information)

In general, all operations of SAP Enterprise Portal that use some type of storage for saving business information must be considered for backup. The backup type to be defined for individual components depends on the type of data storage.

In most cases, backing up just portal databases is sufficient. Data, or information stored in different storage types can be regenerated by the component itself. Therefore, for components that use only database-like storage, database backup is sufficient, but in some installations landscapes, a component can use multiple storage types for storing business information. To address all possible problems, sources of the problem must be identified.

Software and Configuration Data

The system and application software are worth backing up to prevent re-installation in case all, or a part of the software is destroyed. As re-installation is always possible, backup is not mandatory.

Backing up the application at least once after it has been installed, or after it has been upgraded is recommended. Note that restoring a backup of the system, or application is generally only possible if it is restored to exactly the same hardware.

Installing the software often takes time and requires different configurations and customization, thus saving this kind of information may also be important to provide a fast restore and to avoid time and effort after possible failures.

As configurations may change more often than the application itself, the configuration files must be backed up on a regular basis, or each time the configuration is modified.

Instead of identifying all relevant files, and configuration entries, a full system backup will probably be the easiest option for backing up software and configuration data.

The following should be considered when planning a backup and restore strategy for SAP Enterprise Portal, and its configuration files:

- Operating system
- Relational Database Management System (RDBMS) software
- SAP J2EE Engine
- Web server
- Other SAP software and file systems
- Log files (SAP and other)
- Other system components

Provide a backup procedure for Enterprise Portal and its configuration data, such as, providing the specific folders for the SAP J2EE Engine (specified during EP installation), and other components. In some cases, you must exercise caution, as most configuration data for components in the portal are stored in databases. Such data can be stored in the file system,

too. Providing a consistent backup of configuration data in such cases is only possible using off-line mode.

Also note that most components of the portal can regenerate the configuration data again on loading, but special attention must be given to SAP J2EE Engine.



Configuration data for SAP J2EE Engine 6.20 is in the file system. This makes it difficult to implement a consistent online backup solution.

Configuration data for SAP J2EE Engine 6.30 and above is in a database, as such any existing database backup and restore solution can be used.



Identifying Critical Components

Knowing the Enterprise Portal components is helpful in planning your backup strategy. The following are the critical components for the portal.

SAP J2EE Engine 6.20 Cluster	An SAP proprietary Java application server based on the Java 2 Enterprise Edition (J2EE™) standards. The Portal Runtime runs on top of the SAP J2EE Engine.
Portal Server	The Portal Server is logical environment comprised of a collection of software applications for running and managing the portal.
iView Runtime Java (IRJ) application	The iView runtime application is a run time environment based on a Java-servlet for processing iViews. This servlet is in the Portal Server environment.
Portal System Database	This is a relational database management system (RDBMS) used as a repository for the portal. The database is usually on a separate machine. The repository stores data for the Portal Content Directory (PCD), User Management Engine, and Knowledge Management Portal Content Directory (PCD) The PCD is based on the database repository. It stores runtime objects, including role definitions, and page-to-role relationships in the Portal System Database. The PCD also stores deployable (PAR) and iView Templates together with their personalization data and derivations.
User Persistence Store	Refers to user related data stored in one or more repositories. Such user related data repository might be, a database, Lightweight Directory Access Protocol (LDAP) directory server, or SAP R/3 System 6.20.

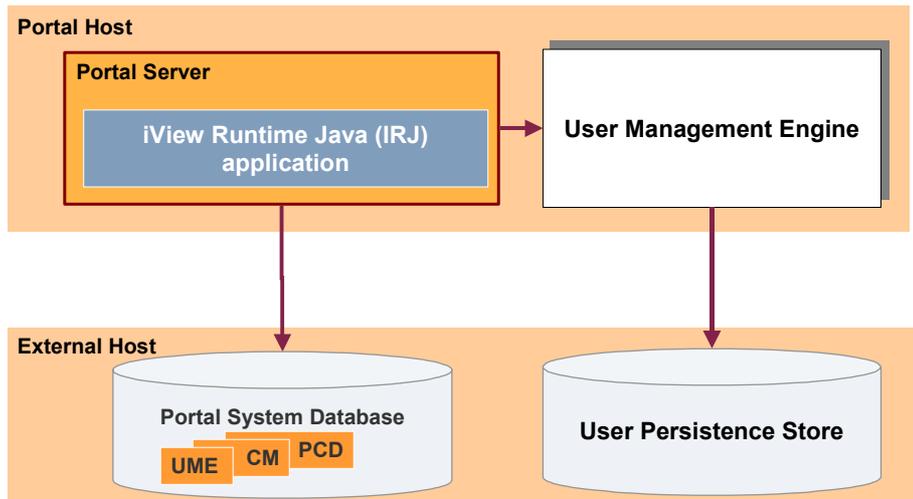
As the portal interrelates with other components, these must be backed up too, such as:

- Web servers
- User Persistence Store
- Java applications and their configuration files
- Native applications and their configuration files
- Relational Database Management System (RDBMS) and their data



If the configuration data is modified during an online backup, TYPE and scope of inconsistencies CANNOT be foreknown.

The following illustration shows the connectivity infrastructure of the critical components:



Identifying the critical components in a backup strategy also depends on the configuration of the deployed Enterprise Portal. Furthermore, knowing the critical components of the portal can be helpful in the implementation of an appropriate backup strategy. Such strategy must take the following into consideration:

- The backup method to implement, and the scheduled interval for performing backup of the component.
- The restoration cost and time of the component.

Relational Database Management System (RDBMS) and their data

The portal related data on a database makes the specific database a critical component. A customer cannot afford to lose any of the user specific data or Enterprise Portal configuration data stored in a database.

LDAP Directory Server Backup

Another critical component is the **LDAP server** and its directory of Enterprise Portal master and configuration data.

Applications and Their Configuration

Depending on the layout of the Enterprise Portal, and the effort that goes into restoring an application, it can be categorized as a critical component.

See also:

[Portal Server \[Page 16\]](#)

[Knowledge Management \[Page 17\]](#)

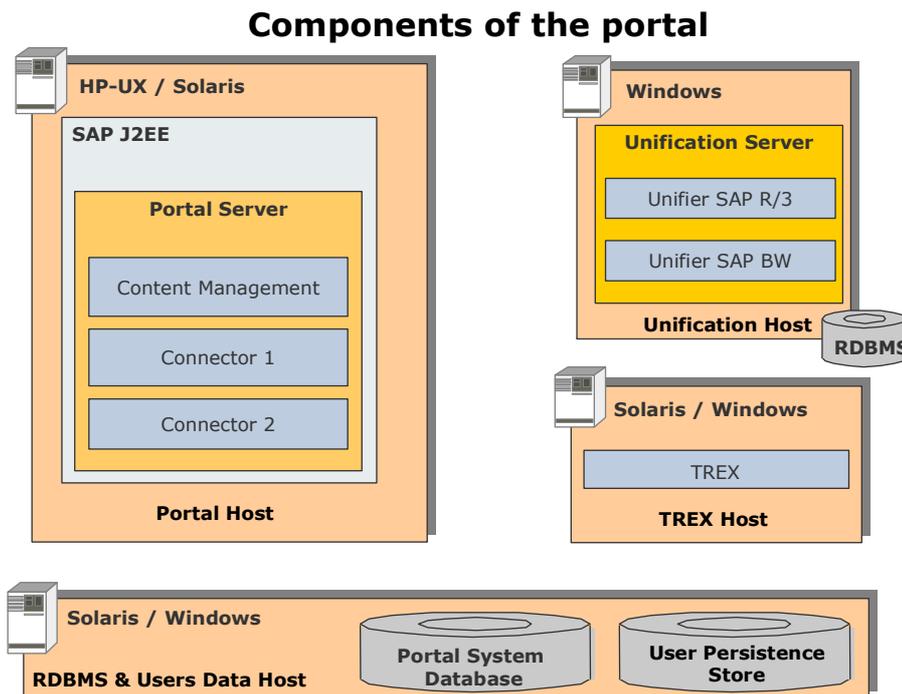
Enterprise Portal Components

SAP Enterprise Portal 6.0 (EP) does not include utilities for performing backup and restore, however, the system administrator can use third-party backup and restore tools.

Within SAP environment, there are existing tools, and backup and restore strategy that be used for some specific components. Such backup can be also used to create a fully functional duplicate of the server at a point-in-time.

A successful backup strategy for SAP Enterprise Portal requires an understanding of how the portal is organized, and configured for the specific customer.

The following is a typical deployment of SAP Enterprise Portal for development and testing:



SAP Enterprise Portal 6.0 runs on the SAP J2EE engine, a proprietary Java application server based on the Java 2 Enterprise Edition (J2EE™) standards. Generally the out-of-the-box portal consists of the following components:

- Portal Server
 - Built on top of the Portal Runtime, the Portal Server components include Portal Content Directory (PCD), Portal System Database (repository database), and User Management Engine.
- Knowledge Management including Content Management, Text Retrieval and Information Exchange (TREX) components.
- Unification components include a Web server, Unification Server, and Unifiers.
- Connectors Framework may include deployed adapters that provide connectivity between an Enterprise Information System (EIS), the SAP J2EE engine, and a specific enterprise application.

See also:

[Identifying the Critical Components \[Page 12\]](#)



Windows

You can implement a backup strategy using the operating system as a basis. On Windows, the following portal components can be identified:

- SAP J2EE Engine
- WEB Server
- Windows NT/2000 services
- Windows NT/2000 applications
- Portal System Database and its data
- User Persistence Store

See also:

[Identifying Critical Components \[Page 12\]](#)



Unix Platform

On Unix, the following critical components can be identified:

- SAP J2EE Engine
- WEB server
- Unix services
- Unix applications
- Portal System Database and its data
- User Persistence Store

See also:

[Identifying Critical Components \[Page 12\]](#)

 **Portal Server**

Portal Server is a logical environment, holding a collection of various software components for running and managing the portal. For instance, among the collection are portal objects that create pages and their defined layouts, manage styles and themes, and provide access to the resources needed at runtime.

Elements used by the runtime components are in the Portal Content Directory (PCD), and these are cached. In this way, a Portal Runtime can gather information from the PCD cache, and obtain the list of iViews related to each page. Data used by Portal Server and the Portal Content Directory (PCD) is stored in Portal System Database.



Documenting your work during implementation of a backup for the first time, is always recommended.

Software Backup

The most comprehensive solution is to implement a full system backup after installation, and configuration. Re-installing the portal is always an option too.

Data Backup

All application and configuration data for Portal Server can be backed up using a standard backup and restore procedure.

See also:

[Knowledge Management \[Page 17\]](#)

[User Management Engine \[Page 20\]](#)



Knowledge Management (KM)

Purpose

The Knowledge Management Platform (KM) provides functions that support the administration of unstructured data, and partly structured data in SAP Enterprise Portal. [Content management \[Page 19\]](#), and retrieval and classification operations are highly integrated. However, retrieval and classification are based on the [Text Retrieval and Information Extraction \(TRES\) \[Page 18\]](#) service. This has its own infrastructure, and can be considered as an individual component.

See also:

[Portal Server \[Page 16\]](#)

[User Management Engine \[Page 20\]](#)



Text Retrieval and Information Extraction (TREC) Software Backup

Purpose

A full system backup after installation and configuration changes is recommended. If you are aware of all configuration entries and changes that have been made on TREC, re-installing the software is an option too.

Data Backup

TREC doesn't have application data in classical sense, because his main functionality is indexing and all information about this processes are stored in index files. Backing up these files can be considered important because time TREC need to re-index all content can be very long.

See also:

[Knowledge Management \[Page 17\]](#)



Content Management (CM)

Purpose

Content management systems are integrated into the Knowledge Management platform, and deployed as a service in the portal.

Software Backup

A full system backup after installation and configuration changes is recommended. If you are aware of all configuration entries and changes that have been made on CM, re-installing the software is an option too.

Data Backup

The CM gives the customers a choice to use different repositories in parallel for storing application data; these have some impact on the backup and restore strategy.

- CM repository in file system - database mode:

In this mode metadata is stored in the DATABASE (properties, locks, etc) and namespace (folder-structure) and content are stored in a file system.

If the database and file system are uncoordinated, the repository is able to update the metadata (the file system is the "reference"). Therefore, no critical issues may arise.

- CM repository in database – file system mode:

The file system is used to store the content of documents, and everything else is stored in the database. In this mode, there can be documents available in the database, but with no content.

- Configuration Data:

There is the likelihood that the configuration data does not match the content of the database. However, since components are added during runtime only (without restart of the servlet engine), there should be no critical issues for a backup.

See also:

[Knowledge Management \[Page 17\]](#)



User Management Engine (UME)

Purpose

SAP User Management Engine (UME) consists of a user management service in the Portal Server that connects to and manages user and group data stored in the User Persistence Store.

User Persistence Store refers to user related data stored in one or more repositories. This repository might be, a database, Lightweight Directory Access Protocol (LDAP) directory server, or SAP R/3 System 6.20.

The UME is interconnected to the User Persistence Store, therefore any backup strategy must consider the type of User Persistence Store implemented at the customer's site.

Scenarios: Backing up User Persistence Store

The following scenarios illustrate backing up user related data in the User Persistence Store configured for the User Management Engine:

- Scenario 1: All users with all attributes in a database
- Scenario 2: All users with basic attributes in LDAP, all other user data in database
- Scenario 3: Some users with basic attributes in LDAP, the other user data in database. Some users completely in database

In Scenario 1, all the information is centrally located in one database, making it easier to implement a backup and restore strategy.

The other scenarios use multiple persistence repositories, where some of the LDAP directory servers do not offer any backup and restore capabilities as done by the common database implementations.

Nevertheless, these scenarios offer other advantages like openness. If a customer chooses to use one of these scenarios (or a similar one), beware of the limitations with such backup and restore strategies.

The following is the recommended practice applicable to all scenarios except scenario 1:

Always backup and restore LDAP and databases at the same time. It is recommended to backup LDAP directly after a database backup. The reason for that is due to the following consequences where only one of LDAP/databases is restored:

- If LDAP is newer compared to database; this happens after restoring only the database, which has been updated, or both are restored and the LDAP backup (time t1) is newer than the database backup (t2).

Consequences:

- User created between t1 and t2

The additional user data is lost for some of the users

- User deleted between t1 and t2

The additional user data exists without users, there might be links to non-existing users in some application tables

- If database is newer compared to LDAP; this happens after restoring only the LDAP, which has been updated, or where both are restored and database backup is newer than the LDAP backup.

Consequences:

- User created between t1 and t2

The additional user data exists without users; there might be links to non-existing users in some application tables.

- User deleted between t1 and t2

The additional user data is lost for some of the users

Software Backup

A full system backup after installation and configuration changes is recommended. If you are aware of all configuration entries and changes that have been made on UME, re-installing the software is an option too.

Data Backup

UME shares the same database schema with Portal Server (PCD), therefore when you backup this database, both application data and configuration data for Portal Server are backed up. In addition, if UME uses LDAP, then you must backup the LDAP database too.

See also:

[Portal Server \[Page 16\]](#)

[Knowledge Management \[Page 17\]](#)

Restoring

A successful backup is good to have, however, a successful restore is the key. Restoring a system from an online backup is always questionable, therefore performing verification tests must follow a restored system based on an online backup.

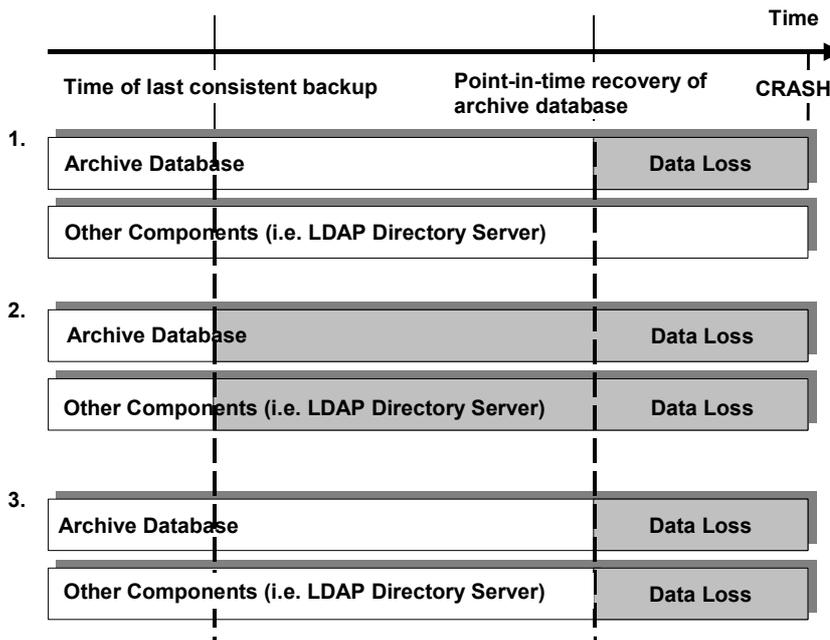
What Can Be Restored?

The answer to this question lies in the type of backup implemented. Since there are no system wide checks, data losses must be taken into account, especially when restoring databases.

Consider a situation where a database holding the archive database has to be restored after a crash. For databases, the database's built-in point-in-time recovery and log recovery can be used. The mechanism allows an almost consistent recovery of the data close to the time of crash.

Databases usually store configuration data for some of the components as well, the applications that store their configuration data in the database would be safe. In addition, these same components store their configuration data on the local file system, so the files would have to be restored to the same point in time, to be consistent. This is not possible using file system backup, and inconsistencies have to be expected. In such situations, you have to weigh the losses.

The following illustration presents likely issues under different scenarios:



In scenario 1, a point in time recovery of the archive database has been implemented. All the other components will keep running. Data loss is then restricted to the archive database only, but the impact of the data loss may be severe.

Scenario 2 shows the complete consistent recovery. All components are affected by the data loss, which might be more severe than in scenario 1.

Scenario 3 shows the situation where a point in time of all affected components was possible. This will be difficult to achieve since some components have been backed up via file system backups. So probably, scenario 3 will not happen for products' landscapes.

The main issue for database recoveries in the future will be the restoration of other dynamic data not stored in a database like the LDAP directory information. That is, some components

also hold dynamic configuration data, which should be restored as close as possible to the point in time of the crash. However, since that data do not change, frequent partial offline backup or online backup is sufficient to get a near consistent state.



Complete tests have not been done, and some of the issues are likely to change, upon testing.



Post - Recovery Checks

After successfully implementing a restore, you must verify whether the restored component works as expected. Two complete checks that can be performed are as follows:

- Can the component be started properly?
You can check if the component is working properly by monitoring the corresponding component's operating system (OS) processes using the operating system's built-in tools, or a professional monitoring tool, or by checking the log files.
- Does the restored and restarted component work properly in the product?
Implementing this check is complicated since you must verify the reconnection of the component to the rest of the product's components (active integration), as well as confirm from the product that the restarted component has been recognized (passive Integration).

See also:

[Restoring \[Page 22\]](#)